arXiv:1407.6128v1 [cs.IR] 23 Jul 2014

Permutation Models for Collaborative Ranking

Truyen Tran and Svetha Venkatesh
Department of Computing, Curtin University, Australia

Feb 2010

Abstract

We study the problem of collaborative filtering where ranking information is available. Focusing on the core of the collaborative ranking process, the user and their community, we propose new models for representation of the underlying permutations and prediction of ranks. The first approach is based on the assumption that the user makes successive choice of items in a stage-wise manner. In particular, we extend the Plackett-Luce model in two ways - introducing parameter factoring to account for user-specific contribution, and modelling the latent community in a generative setting. The second approach relies on log-linear parameterisation, which relaxes the discrete-choice assumption, but makes learning and inference much more involved. We propose MCMC-based learning and inference methods and derive linear-time prediction algorithms.

Keywords: permutation, ranking, collaborative filtering.

1 Introduction

Collaborative filtering is an important class of problems with the promise to deliver personalised services. Members of communities rate items in a service, and strong patterns exist between similar communities of users. These patterns can be exploited to produce ranked lists of items from a set of items not not previously exposed to the user.

Research in recommendation systems models user preferences through a numerical rating - for example, rate a movie as 4 or 5 stars. Although these users are forced into numeric scoring, these scores are assigned qualitatively, and do not carry the assumed rigour of quantitative evaluation. Also, this limits the expressiveness of preferences. For example, a more intuitive way is to express the order of preferences for a set of items. It may be easier to rank a set of movies, or the top 10 places visited, rather than assign them a numeric score. Importantly, in recommendation systems, the core value proposition is to recommend unseen items - this is where ranking rather than actual rating becomes significant.

This paper addresses the open problem of recommending a ranked list of items, or a preference list, without requiring intermediate ratings, in collaborative filtering systems. Each user provides a ranked list of items, in the decreasing order of preference. The list needs not be complete, e.g. a user typically rates 10 or 20 items. The intuition in collaborative filtering is that the community as a whole may cover thousands of items,

and as users belong to clusters within this community, the properties of rankings within such clusters can be transferred to a user for items that user has not seen. The technical issue is to model the ranked item set both for a user and the community, and predict the rank of unseen items for each user.

Despite of its importance, the collaborative ranking problem has only been attempted recently [10, 11, 7]. The papers [11, 7] consider pairwise preferences , ignoring the simultaneous interaction between items. Listwise approaches, studied in statistics (e.g. see [8, 9, 4]), often involve a relatively small set of items (e.g. in election, typically less than a dozen of candidates are considered). Further, statisticians are interested in the distribution of ranks in the population rather than properties of individuals. Collaborative ranking, on the other hand, differs in three ways: a) the scale is significantly different - sometimes there are millions of items; b) the data is highly sparse, that is a user will typically users will only express their preference over a few items; and c) the personalisation aspect is crucial and this the the distribution of rank per user is more important.

In this paper, focusing on the user, we study two approaches for modelling the rank or preference lists. Our first approach assumes that the user, when ranking items, will make successive choice in a stage-wise matter. We extend one of the most well-known methods, namely the Plackett-Luce model, to effectively model user-specific rank distribution in two ways. First, we introduce parameter factoring into user-specific and item-specific parameters. Second, we employ a generative framework which models the community the user belongs to as a latent layer, enabling richer modelling of the community structure in the ranking generation process. We provide algorithms for learning the model parameters and for ranking unseen items in linear time. The approach is detailed in Section 3.

The second approach relaxes the stage-wise choice assumption and models intrinsic features of the permutation in a log-linear setting. Potentials in the model capture the likelihood of an item in a specific position, and for all item pairs, the likelihood of the first item being ordered before the second. Although exact learning and inference is intractable, we show that truncated MCMC techniques are effective for learning, and for prediction that can be computed in linear time. The approach is described in Section 4.

The novelty in our contribution lies in the proposal of two approaches incorporating key aspects of collaborative ranking: the user, their specific communities, and the nature of the ranking list itself. The work contributes efficient methods for learning and prediction.

2 Preliminaries

Suppose that we have a data set of N users, and M items and each user $u \in \{1, 2, ..., N\}$ provides a list of $n_u \leq M$ ranked items $\pi^u = \{\pi^u_1, \pi^u_2, ...\pi^u_{n_u}\}$, where π^u_i is the index of the item in position i. For notational simplicity, we will drop the explicit superscript u in π^u when there is no confusion, and use $y = \pi_i$ when we mention the item $y \in \{1, 2, ..., M\}$ in position i. The goal is to effectively model the distribution $P(\pi|u)$. The main difficulty is that the number of permutations is $n_u!$, which is only tractable for small n_u.

A simplified way is to examine the ordering between only two items (e.g. see [11, 7]). Denote by $s^u_{\pi_i}$ the scoring function when the item is positioned at i in the list π of user u. Let us consider the following quantity

$$d^u_{ij} \quad = \quad \text{sign}(j - i)(s^u_{\pi_i} - s^u_{\pi_j}) \, .$$

Basically d_{ij}^u is positive when the scoring functions $\{s_y^u, s_{y'}^u\}$ agree with their relative positions in the list, and negative otherwise. For simplicity, we assume the factoring $s_y^u = \sum_{k=1}^K W_{uk} H_{ky}$ where $W \in \mathbb{R}^{N \times K}$ and $H \in \mathbb{R}^{K \times M}$ for some $K < \min\{M, N\}$. Thus the learning goal is to estimate $\{W, H\}$ so that $\{d_{ij}^u\}$ are positive for all the triples (u, i, j) in the training data, where $1 \leq i < j \leq n_u$. This suggests a regularised loss function in the form

$$\mathcal{R} = \frac{1}{N} \sum_u \sum_{i=1}^{n_u} \sum_{j=i+1}^{n_u} L(d_{ij}^u) + \Omega(W, H)\,,$$

where $L(d_{ij}^u)$ is the user-specific loss and $\Omega(W, H) = \alpha \sum_{uk} W_{uk}^2 + \beta \sum_{yk} H_{ky}^2$ is the regularising component. Popular choices of $L(d_{ij}^u)$ are

$$L(d_{ij}^u) = \begin{cases} (1 - d_{ij}^u)^2 & \text{in regression;} \\ \max(0, 1 - d_{ij}^u) & \text{in large-margin setting; and} \\ \log(1 + \exp\{-d_{ij}^u\}) & \text{in logistic regression.} \end{cases}$$

3 Latent Discrete Choice Models

We now address the listwise models, starting from the assumption that the user makes the ranking decision in a stage-wise manner. We will focus on the Plackett-Luce model [9]

$$P(\pi) = \prod_{i=1}^{M} \frac{e^{s_{\pi_i}}}{\sum_{j=i}^{M} e^{s_{\pi_j}}}\,, \tag{1}$$

where s_{π_i} is the score associated with the item at position i in the permutation π. The probability that an item is chosen as the first in the list is $e^{s_{\pi_1}} / \sum_{j=1}^{M} e^{s_{\pi_j}}$. Once this item has been chosen, the probability that the next item is chosen as the second in the remaining of $M - 1$ item list is $e^{s_{\pi_2}} / \sum_{j=2}^{M} e^{s_{\pi_j}}$. The process repeats until all items have been chosen in appropriate positions.

However, this model is not suitable for collaborative ranking, because it does not carry any personalised information, and lacks the concept of community among users. We now introduce our extensions, first by modelling the user-specific distribution $P(\pi|u)$ (Section 3.1), and then proposing community-generated choice making (Section 3.2).

3.1 Factored Benter-Plackett-Luce Model

In collaborative ranking, we are interested in modelling the choices by each user, and the permutation π given by a user is incomplete (i.e. the user often ranks a very small subset of items). We thus introduce an user-specific model as

$$P(\pi|u) \quad = \quad \prod_{i=1}^{n_u} \frac{e^{s_{\pi_i}^u}}{\sum_{j=i}^{n_u} e^{s_{\pi_j}^u}}\,.$$

Thus $s_{\pi_i}^u$ is the ranking score for item at position i (under π) by user u. However, this model does not account for the the order at the beginning of the list being more

important than that at the end. We employ the technique by [1], introducing *damping* factors $\rho_1 \geq \rho_2 \geq ... \geq \rho_n \geq 0$ as follows

$$P(\pi|u) = \prod_{i=1}^{n_u} \frac{e^{\rho_i s^u_{\pi_i}}}{\sum_{j=i}^{n} e^{\rho_i s^u_{\pi_j}}} .$$

As an example, we may choose $\rho_i = 1/\log(1+i)$.

In the standard Plackett-Luce model, the set of parameters $\{s_y\}$ can be estimated from a set of $i.i.d$ permutation samples. In our adaptation, however, this trick does not work because the score s^u_y will be undefined for unseen items. Instead, we propose to factor s^u_y as follows

$$s^u_y = \sum_{k=1}^{K} W_{uk} H_{ky} ,$$

where $W \in \mathbb{R}^{N\times K}$ and $H \in \mathbb{R}^{K\times M}$ for some $K < \min\{M, N\}$ are parameter matrices. The yth column of H can be considered as the feature vector of item y, and the uth row of W as the parameter vector specific to user u.

To learn the model parameters, maximum likelihood estimation can be carried out through maximising the following regularised log-likelihood with respect to $\{W, H\}$

$$\mathcal{L}(W, H) = \sum_u \log P(\pi|u) - \alpha \|W\|^2_F - \beta \|H\|^2_F ,$$

for $\alpha, \beta > 0$. It can be verified that the regularised log-likelihood is concave in either W or H, but not both. Once the model has been specified, $\{s^u_y = \sum_{k=1}^{K} W_{uk} H_{ky}\}$ can be used for sorting the items previously not seen by the user, where larger s^u_y ranks the item higher in the list.

3.2 Latent Semantic Plackett-Luce Model

The model in the previous subsection lacks generative interpretation- we do not know how the ranking is generated by the user. A principled way is to assume that the user belongs to hidden communities, and that those communities will jointly generate the ranking. Recall that in the Plackett-Luce model, the choice of items is made stage-wise - the next item is chosen given that previously chosen items are ahead in the list. Denote by $P_i(\pi|z, u)$ the probability of choosing the item for the ith position by u with respect to community z, i.e.

$$P_i(\pi|z, u) = \frac{e^{s^z_{\pi_i}}}{\sum_{j=i}^{n_u} e^{s^z_{\pi_j}}} . \tag{2}$$

Let $P(z|u)$ be the probability that the user belongs to one of the communities $z \in \{1, 2, .., K\}$, then the user-specific permutation is defined as

$$P(\pi|u) \quad = \quad \prod_{i=1}^{n_u} \sum_z P(z|u) P_i(\pi|z, u) . \tag{3}$$

Due to the sum in the denominator in Equation 2, we may expect that the computation of $P(\pi|u)$ takes $n_u(n_u-1)K/2$ time. However, we can compute in n_uK time by precomputing a recursive array $A_i^z = A_{i+1}^z + e^{s^z_{\pi_i}}$ for $1 \le i < n_u$. If we start with $A_{n_u} = e^{s^z_{\pi_{n_u}}}$, then clearly $A_i^z = \sum_{j=i}^{n_u} e^{s^z_{\pi_j}}$, which is the denominator in Equation 2.

3.2.1 Learning using EM

There are two sets of parameters to estimate, the mixture coefficients $\{P(z|u)\}$ and the community-specific item scores $\{s_y^z\}$. We describe an EM algorithm for learning these parameters, starting from the lower-bound of the incomplete log-likelihood $\mathcal{L} = \sum_u \log P(\pi|u)$ as

$$\begin{aligned}\mathcal{L} &= \sum_u \sum_{i=1}^{n_u} \log \sum_z P(z|u) P_i(\pi|z,u) \\ &\ge \sum_u \sum_{i=1}^{n_u} \sum_z Q_i(z|\pi,u) \log P(z|u) P_i(\pi|z,u) \\ &= \mathcal{Q}\,.\end{aligned}$$

where $Q_i(z|\pi,u)$ is defined at each **E**-step $t+1$ as follows

$$Q_i^{t+1}(z|\pi,u) \leftarrow \frac{P^t(z|u)P_i^t(\pi|z,u)}{P_i^t(\pi|u)}\,.$$

In the **M**-step, we fix $Q_i(z|\pi,u)$ and estimate $\{P(z|u),\, s_y^z\}$ by maximising $\mathcal{Q}$. We equip the lower-bound with the constraint $\sum_z P(z|u) = 1$ through the Lagrangian function $\mathcal{F} = \mathcal{Q} + \sum_u \mu_u(\sum_z P(z|u) - 1)$ where $\{\mu_u\}$ are Lagrange multipliers. Setting the gradient of the Lagrangian function

$$\frac{\partial \mathcal{F}}{\partial P(z|u)} = \sum_{i=1}^{n_u} Q_i(z|\pi,u)\frac{1}{P(z|u)} + \mu_u$$

to zeros and maintaining that $\sum_z P(z|u) = 1$ would lead to

$$\begin{aligned}P(z|u) &\leftarrow \frac{\sum_{i=1}^{n_u} Q_i(z|\pi,u)}{\sum_z \sum_{i=1}^{n_u} Q_i(z|\pi,u)} \\ &= \frac{1}{n_u}\sum_{i=1}^{n_u} Q_i(z|\pi,u)\,.\end{aligned}$$

This closed form update, however, does not apply to $\{s_y^z\}$. Instead, we resort to the gradient-based method, where

$$\begin{aligned}\frac{\partial \mathcal{Q}}{\partial s_y^z} &= \sum_{i=1}^{n_u} Q_i(z|\pi,u)\frac{\partial \log P_i(\pi|z,u)}{\partial s_y^z} \\ &= \sum_{i=1}^{n_u} Q_i(z|\pi,u)\{\delta^y_{\pi_i} - \frac{\sum_{j=i}^{n_u} e^{s^z_{\pi_j}}\delta^y_{\pi_j}}{\sum_{j=i}^{n_u} e^{s^z_{\pi_j}}}\}\,,\end{aligned}$$

where $\delta^y_{\pi_i} = 1$ if $y = \pi_i$ and 0 otherwise. Typically, we run only a few updates for s^z_y per **M**-step.

3.2.2 Prediction

Given that models are fully specified, we want to output a ranked list of unseen items for each user u. However, finding the optimal ranking for an arbitrary set of items is generally intractable and thus we resort to finding the rank of just one unseen item at a time, given that the seen items have been sorted. In other words, we fix the orders of the old items, and then introduce one new item into the model, assuming that this introduction does not change the relative orders of the old items. So the problem now reduces to finding the position of the new item among the old items.

We repeat the process for all new items, and determine their positions in the list. If the two new items are placed in the same position, then their relative ranks will be determined by the likelihood of their introductions.

Let π' be the new list after introducing a new item. Denote by $\pi_{i:j}$ the set of items whose positions are from i to j under π. Suppose that the new item is placed between the $(j-1)$th and the jth items of the the old list π, and thus it is in the jth position of the new list π'. Thus $\pi'_{1:j-1} = \pi_{1:j-1}$ and $\pi'_{j+1:n+1} = \pi_{j:n}$. We want to find

$$j^* = \arg\max_j P(\pi'_{1:j-1}, \pi'_j, \pi'_{j+1:n+1}|u)\,,$$

where $P(\pi'_{1:j-1}, \pi'_j, \pi'_{j+1:n+1}|u) =$

$$\left[\prod_{i=1}^{j-1}\sum_z P(z|u)P_i(\pi'|z,u)\right]\left[\sum_z P(z|u)P_j(\pi'|z,u)\right] \quad \times$$
$$\times \left[\prod_{i=j+1}^{n+1}\sum_z P(z|u)P_i(\pi'|z,u)\right].$$

Naive computation for finding the optimal j^* will cost $n_u(n_u+1)K/2$ steps. Here we provide a solution with just $(n_u+1)K$ steps. We will proceed from left-to-right in a recursive manner, starting from $j = 1$. Recall that we can compute $P(\pi'_{1:n+1}|u)$ in Equation 3 in $(n_u+1)K$ steps.

Assume that we have computed for the case that the position of the new item is j (under π'), we want to compute the case that the new position is $j+1$ (under π''). Let us examine the odds

$$O_j = \frac{P(\pi''_{1:j}, \pi''_{j+1}, \pi''_{j+2:n+1}|u)}{P(\pi'_{1:j-1}, \pi'_j, \pi'_{j+1:n+1}|u)}\,.$$

We have $P(\pi''_{1:j}, \pi''_{j+1}, \pi''_{j+2:n+1}|u) =$

$$\left[\prod_{i=1}^{j-1}\sum_z P(z|u)P_i(\pi''|z,u)\right]\left[\sum_z P(z|u)P_j(\pi''|z,u)\right] \times$$
$$\times \left[\sum_z P(z|u)P_{j+1}(\pi''|z,u)\right]\left[\prod_{i=j+2}^{n+1}\sum_z P(z|u)P_i(\pi''|z,u)\right].$$

We now notice that $\pi''_{1:j-1} = \pi'_{1:j-1}$ and $\pi''_{j+2:n+1} = \pi'_{j+2:n+1}$, and

$$\begin{aligned} P_i(\pi'|z) &= P_i(\pi''|z) \\ \forall z, \text{ and for } i &\in \{1:j-1\} \cup \{j+2:n_u+1\}. \end{aligned}$$

The odds can be simplified as

$$O_j = \frac{\left[\sum_z P(z|u) P_j(\pi''|z,u)\right]\left[\sum_z P(z|u) P_{j+1}(\pi''|z,u)\right]}{\left[\sum_z P(z|u) P_j(\pi'|z,u)\right]\left[\sum_z P(z|u) P_{j+1}(\pi'|z,u)\right]}. \tag{4}$$

which costs K time to evaluate. Consequently, the recursive process costs totally $(n_u + 1)K$ time steps.

4 Log-linear Models

In this section, we propose a second approach to permutation modelling. The main difference from the first approach is that we do not make the discrete-choice assumption, which makes the parameter estimation easy, but complicates the inference. We now rely on the log-linear parameterisation, which is more flexible. The generic conditional distribution is defined as

$$P(\pi|u) = \frac{1}{Z(u)} \left[\prod_{i=1}^{n_u} \phi_\pi(i,u)\right] \left[\prod_{i=1}^{n_u-1} \prod_{j=i+1}^{n_u} \phi_\pi(i,j)\right], \tag{5}$$

where $\phi_\pi(i,u)$ and $\phi_\pi(i,j)$ are positive potential functions, $Z(u)$ is the normalising constant (a.k.a partition function). The position-wise potential $\phi_\pi(i,u)$ captures the likelihood that a particular item $y = \pi_i$ is placed in position i by user u. For example, we would expect that a particular movie is among the top 5% in the list of a user. On the other hand, the pairwise potential $\phi_\pi(i,j)$ encodes the likelihood that the item $y = \pi_i$ is preferred to item $y' = \pi_j$. In what follows, we will make use of the energy notion, i.e. $\phi_\pi(i,u) = \exp\{-E(\pi_i,u)\}$ and $\phi_\pi(i,j) = \exp\{-E(\pi_i,\pi_j)\}$. The energy of the permutation π is therefore the sum of component energies, i.e. $E(\pi,u) = \sum_i E(\pi_i,u) + \sum_i \sum_{j>i} E(\pi_i,\pi_j)$.

4.1 MCMC for Inference

Inference in the above generic model is intractable due to the partition function $Z(u)$, which requires $\frac{1}{2}n_u^2(n_u-1)^2(n_u-2)!$ computational steps[1]. We thus resort to MCMC methods. The key is to design a proposal distribution that helps the random walks to quickly reach the high density regions. There is also a trade-off here because large steps would mean significant distortion of the current permutation, resulting in more computational cost per move. We consider three types of local moves.

Item relocation. Randomly pick one item in the list, and relocate it, keeping the relative orders of the rest unchanged. For example, assume the permutation is $[A,B,C,D,E,F]$ and if B is relocated to the place between E and F, then the new permutation is $[A,C,D,E,B,F]$. Generally, this type of move costs $\mathcal{O}(n_u)$ operations per move due

[1] There are $n_u!$ permutations, each require $\frac{1}{2}n_u(n_u-1)$ steps of computing the product of potentials.

to the change in relative preference orders. In the example we are considering, the pairs BC, BD, BE would change to CB, DB, EB.

Item swapping. Randomly pick two items, and swap their positions leaving other items unchanged. In the above example, if we swap B and E, then the new permutation is $[A, E, C, D, B, F]$. This also costs $\mathcal{O}(n_u)$ operations per move.

Sublist permutation. Randomly pick a small sublist, try all permutations within this sublist. For example, the sublist $[B, C, D]$ will result in $[C, B, D], [B, D, C], [D, C, B], [C, D, B], [D, B, C]$. This costs $\Delta!$ where Δ is the size of the sublist. When $\Delta = 2$, this is the special case of the item swapping.

Since the proposals are symmetric, the acceptance probability in the Metropolis-Hastings method is simply

$$P = \min\{1, e^{-\Delta E}\}\,, \tag{6}$$

where ΔE is the change in model energy due to the proposed move.

4.2 Learning with Truncated MCMC

Learning using maximum likelihood is intractable due to the computation of $Z(u)$ and its gradient, and thus MCMC-based learning can be employed. The assumption is that if we generate enough samples according to the model distribution, then the gradient of the log-likelihood can be accurately estimated, and thus learning can proceed. However, this is clearly too expensive, because generally we would need a significantly large number of samples per gradient evaluation. Instead, Hinton [5] proposes a simple technique called Contrastive Divergence (CD) that has been shown to work well in standard Boltzmann machines. The idea is that instead of starting the Markov chain randomly and running forever, we can just start from the observed configuration, and run for a few steps. This is enough to relax the model from the empirical distribution.

Here we adopt the CD, but we should stress in passing that the application of CD in the context of permutation modelling is novel. It is possible that we just need to run one short Markov chain of length $n_u$with the item-swapping moves.

4.3 Learning with Pseudo-likelihood

In standard graphical models, pseudo-likelihood is an efficient alternative to the full likelihood, and it is provably consistent given sufficient regularity in the model structure. However, this concept has no straightforward application in permutation models. We attempt to consider the pseudo-likelihood concept from a more abstract level.

There is a close relationship between pseudo-likelihood and MCMC techniques. The difference is that in MCMC we randomly choose one local permutation configuration, while in pseudo-likelihood, we consider all local configurations, and thus the process is deterministic. Using this idea, the (log) pseudo-likelihood can be written as

$$\begin{aligned} \mathcal{L}^{pseudo} &= \sum_u \sum_c \log P(\pi_c | \pi_{\neg c}, u) \text{ where} \\ P(\pi_c | \pi_{\neg c}, u) &= \frac{\exp\{-E(\pi_c, \pi_{\neg c}, u)\}}{\sum_{\pi'_c} \exp\{-E(\pi'_c, \pi'_{\neg c}, u)\}}\,. \end{aligned}$$

and c denotes the index of the local structure, and $\neg c$ denotes the rest of the items whose relative positions remain unchanged. We briefly discuss three types of local structure.

Item relocation. All the items will be considered, each has the following local distribution

$$P(\pi_i|\pi_{\neg i}, u) = \frac{\exp\{-E(\pi_{1:i-1}, \pi_i, \pi_{i+1:n}, u)\}}{\sum_{j=1}^{n} \exp\{-E(\pi'_{1:j-1}, \pi'_j, \pi'_{j+1:n}, u)\}}$$

for $1 \leq i \leq n_u$. Since the denominator is the sum over n_u positions, each requires $n_u - 1$ pairwise energies, naively computing $P(\pi_i|\pi_{\neg i}, u)$ would result in $n_u(n_u - 1)$ steps. However, we can the the denominator in a single pass. Suppose the item $y = \pi_i$ moves from current position j (under π') to $j + 1$ (under π''), then change in energy is

$$\Delta E_j(\pi' \to \pi'', u) = E(\pi''_j, \pi''_{j+1}, u) - E(\pi'_j, \pi'_{j+1}, u)\,,$$

which costs a constant time to compute. We can start with $j = 1$, updating model energies in one pass.

Item swapping. We have $n_u(n_u - 1)/2$ item pairs for each user u. So the local distribution is

$$P(\pi_{i,j}|\pi_{\neg i,j}, u) = \frac{1}{1 + \exp\{-\Delta E_{ij}(u)\}}$$

for $1 \leq i < j \leq n_u$ where $\Delta E_{ij}(u)$ is the change in energy as a result of the swapping items $y = \pi_i$ and $y' = \pi_j$.

Sublist permutation. We will have $n_u + 1 - \Delta$ local distributions of the following form $P(\pi_{i:i+\Delta-1}|\pi_{\neg i:i+\Delta-1}, u) =$

$$\frac{\exp\{-E(\pi_{1:i-1}, \pi_{i:i+\Delta-1}, \pi_{i+\Delta:n}, u)\}}{\sum_{\pi'_{j:j+\Delta-1}} \exp\{-E(\pi'_{1:j-1}, \pi'_{j:j+\Delta-1}, \pi'_{j+\Delta:n}, u)\}}$$

for $1 \leq i \leq n_u + 1 - \Delta$.

4.4 Prediction

We employ the same technique described earlier with the Latent Plackett-Luce model (Section 3.2.2) in that we fix the relative order of the items the user has already seen, and introduce the new item into the list. Then we search for the best position of the new item in the list, where the best position has the lowest permutation energy. Computationally, this is similar to the pseudo-likelihood with item-relocation, except that now we choose the most probable position instead of summing over all positions. Thus, we can find the best position in a single pass.

4.5 Parameterisation Case Studies

We now specify the parameters for the log-linear modelling. We will focus on two special cases, one with factored position-wise parameters, and the other with pairwise parameters.

4.5.1 Factored Position-wise Parameters

Let us start from the idea of augmenting each item with a score s^u_y, which we assume the factored form as $s^u_y = \sum_{k=1}^{K} W_{uk} H_{ky}$. Ignoring the pairwise potentials in Equation 5, the position-wise potential can be defined as $\phi_\pi(i, u) = \exp\{s^u_{\pi_i} g(i, u)\}$ where $g(i, u)$ is

a monotonically *decreasing* function in i. This case is attractive because a MCMC step with position swapping costs only a constant time, i.e. if we swap two items at positions l and m, the change in energy is $\Delta E_{lm}(u) = 2(s^u_{\pi_l} - s^u_{\pi_m})(m-l)$. In addition, prediction is rather simple as we just need to use s^u_y for sorting.

In particular, we are interested in the case $g(i,u) = (1 + n_u - 2i)/n_u$ since it has a nice interpretation

$$\begin{aligned} P(\pi|u) &= \frac{1}{Z(u)} \exp\{\frac{1}{n_u} \sum_{i=1}^{n_u} s^u_{\pi_i}(1 + n_u - 2i)\} \\ &= \frac{1}{Z(u)} \exp\{\frac{1}{n_u} \sum_{i=1}^{n_u-1} \sum_{j=i+1}^{n_u} (s^u_{\pi_i} - s^u_{\pi_j})\}, \end{aligned}$$

which basically says that when $y = \pi_i$ is preferred to $y' = \pi_j$, then we should have $s^u_y > s^u_{y'}$.

4.5.2 Pairwise Parameters

We now consider the second special case, where the pairwise potential is simply $\phi_\pi(i,j) = \exp\{\lambda_{yy'}\}$ subject to $y = \pi_i$ and $y' = \pi_j$. Note that $\lambda_{yy'} \neq \lambda_{y'y}$. Since the total parameters can be as much as M^2, which is often too large for robust estimation, we keep only the parameters of the item pairs whose number of co-occurrences in the training data is larger than a certain threshold. To account for missing pairs, we also use the position-wise potential $\phi_\pi(i,u) = \exp\{\gamma_{\pi_i} g(i,u)\}$ with an extra parameter per item γ_y (here $y = \pi_i$). The distribution is now defined as

$$P(\pi|u) = \frac{1}{Z(u)} \exp\left\{ \sum_{i=1}^{n_u} \gamma_{\pi_i} g(i,u) + \sum_{i=1}^{n_u-1} \sum_{j=i+1}^{n_u} \lambda_{\pi_i,\pi_j} \right\}.$$

For example, the threshold may be set to 5 and we can use $g(i,u) = 1 - i/n_u$. Note that there is no user-specific parameter. However, the distribution is still user-dependent because of the number of items n_u and the ranking are user-specific.

In MCMC, suppose we swap items at positions l and m, where $l < m$, the change in energy is

$$\begin{aligned} \Delta E_{lm}(u) &= (\gamma_{\pi_l} - \gamma_{\pi_m})\{g(l,u) - g(m,u)\} + \\ &\quad \lambda_{lm} - \lambda_{ml} + \sum_{l<i<m} \{\lambda_{li} + \lambda_{im} - \lambda_{il} - \lambda_{mi}\}. \end{aligned}$$

5 Related Work

Although collaborative filtering with numerical ratings is well studied, collaborative ranking is more recent. Work of [10] introduces CoFi^{RANK} - a non-probabilistic method which optimises the bound of the NDCG score. The authors discuss several pairwise loss functions mentioned in Section 2. An adaptation of PLSA [6] for pairwise preference is given in [7]. None of these papers attempt to model the rank distributions.

In statistics, on the other hand, rank models are well-studied, some of which we have already mentioned in previous sections: the Bradley-Terry [2] for pairwise preferences, the Plackett-Luce [9] for discrete choices, the Mallows for rank aggregation [8], and the spectral decomposition [4]. Statistical data, however, is often limited to small sets of items (e.g. less than a dozen) and the goal is to model a single distribution for all users. Our work, on the other hand, deals with large sets of sparsely ranked items, and models user-specific distributions.

Rank learning has recently attracted much attention in Information Retrieval. For example, the Plackett-Luce model has been adapted in [3]. However, the setting is different, since the items (e.g. documents or images) are associated with pre-computed features, and the parameters are only associated with these features, not the items. Collaborative ranking, on the other hand, discovers these features directly from the data.

6 Conclusion

We have studied two approaches of permutation modelling for collaborative ranking under different assumptions. The first approach follows the Plackett-Luce's discrete-choice assumption. We introduce parameter factoring as well as latent semantic extensions to account for hidden community structure among users. The second approach relies on log-linear parameterisation. We show how to perform MCMC-based inference, learning, and efficient recommendation. Future directions include extensions to deal with ties among ranks, and to incorporate correlation between users.

References

[1] W. Benter. Computer based horse race handicapping and wagering systems: A report. *Efficiency of Racetrack Betting Markets (2008 Edition)*, 1994.

[2] R.A. Bradley and M.E. Terry. Rank analysis of incomplete block designs. *Biometrika*, 39:324–345, 1952.

[3] Z. Cao, T. Qin, T.Y. Liu, M.F. Tsai, and H. Li. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, page 136. ACM, 2007.

[4] P. Diaconis. A generalization of spectral analysis with application to ranked data. *The Annals of Statistics*, pages 949–979, 1989.

[5] G.E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14:1771–1800, 2002.

[6] T. Hofmann. Unsupervised learning by probabilistic latent semantic analysis. *Machine Learning*, 42(1):177–196, 2001.

[7] N.N. Liu, M. Zhao, and Q. Yang. Probabilistic latent preference analysis for collaborative filtering. In *CIKM*, pages 759–766. ACM, 2009.

[8] C.L. Mallows. Non-null ranking models. I. *Biometrika*, 44(1):114–130, 1957.

[9] R.L. Plackett. The analysis of permutations. *Applied Statistics*, pages 193–202, 1975.

[10] M. Weimer, A. Karatzoglou, Q. Le, and A. Smola. CoFi^{RANK}-maximum margin matrix factorization for collaborative ranking. *Advances in neural information processing systems*, 20:1593–1600, 2008.

[11] Markus Weimer, Alexandros Karatzoglou, and Alexander J. Smola. Adaptive collaborative filtering. In *Proceedings of the 2008 ACM conference on Recommender systems*, pages 275–282, Lausanne, Switzerland, 2008.